On Scaling Rules for Energy of VLSI Polar Encoders and Decoders

Christopher G. Blake and Frank R. Kschischang
Department of Electrical & Computer Engineering
University of Toronto
christopher.blake@mail.utoronto.ca frank@comm.utoronto.ca

arXiv:1602.04034v1 [cs.IT] 12 Feb 2016

Abstract

It is shown that all polar encoding schemes of rate $R > \frac{1}{2}$ of block length N implemented according to the Thompson VLSI model must take energy $E \geq \Omega\left(N^{3/2}\right)$. This lower bound is achievable up to polylogarithmic factors using a mesh network topology defined by Thompson and the encoding algorithm defined by Arikan. A general class of circuits that compute successive cancellation decoding adapted from Arikan's butterfly network algorithm is defined. It is shown that such decoders implemented on a rectangle grid for codes of rate $R > 2/3$ must take energy $E \geq \Omega(N^{3/2})$, and this can also be reached up to polylogarithmic factors using a mesh network. Capacity approaching sequences of energy optimal polar encoders and decoders, as a function of reciprocal gap to capacity $\chi = (1 - R/C)^{-1}$, have energy that scales as $\Omega\left(\chi^{5.325}\right) \leq E \leq O\left(\chi^{7.05} \log^4(\chi)\right)$.

I. Introduction

CONTINUING recent work on the energy of VLSI error control circuits [1]–[6], in this paper we provide lower and upper bound scaling rules on the energy of VLSI implementations of polar encoders and decoders. In particular, we show that all polar encoders of constant switching activity factor for codes of rate $R > 1/2$ have energy that scales at least as $\Omega\left(N^{3/2}\right)$. We describe a class of circuits based on the polar decoding algorithm suggested by Arikan in [7]. We show that circuits of this type for polar codes of rate $R > 2/3$ must take at least $\Omega(N^{3/2})$ energy when its output nodes are arranged on a rectangular grid. The mesh network topology can also reach this lower bound up to polylogarithmic factors by using circuits with switching activity factor that decreases with increasing block length.

In a companion paper [5], we analyze the tradeoffs in the number of clock cycle, energy, and block error probability for encoder and decoder circuits. For sequences of circuits with variable switching activity factor, Grover [1] showed similar scaling rules for the energy of encoders and decoders as a function of block error probability. In particular, these results show that coding schemes with block error probability that scales exponentially in block length N must have energy that scales as $\Omega(N^{3/2})$, and our paper [5] shows that schemes with constant switching activity factor that reach this lower bound must have number of clock cycles $T \geq \Omega(\sqrt{N})$. There exist generalized polar decoders with asymptotic block error probability that scale as $\Theta(e^{-N^{1-\epsilon}})$ for any $\epsilon > 0$ [8] (that is, very close to $O(e^{-N})$), and in this paper we discuss how the energy of polar decoders for such codes implemented on a mesh network can get very close to the $\Omega(N^{3/2})$ energy lower bound implied by Grover [1], implying that the energy lower bound is close to tight. However, this requires decoders with switching activity factor that scales as $\Theta(1/N)$, and number of clock cycles that scales close to $\Theta(N^{3/2})$, in contrast to the clock cycle lower bound of [5]. This is because of the difficulty of parallelization of the successive cancellation decoding algorithm.

In Section II we discuss how the results of this paper build upon prior work. In Section III we discuss the computational model we will use and present some basic definitions that will be used. In Section IV we present one of the main technical results of this paper, showing a lower bound on the VLSI energy complexity of polar encoding. We discover a similar lower bound for the complexity of decoding using VLSI circuits derived from Arikan's successive cancellation decoding algorithm in Section V. Upper bounds that reach the lower bounds of the previous two sections are presented in Section VI where the mesh network used by Thompson for sorting and fast Fourier transform is applied to the polar encoding and decoding algorithms. In Section VII we study how some of these results can be extended to polar coding with more general generating matrices. In Section VIII we show how these upper and lower bound results, when combined with finite length analysis of polar coding, results in upper and lower bounds for the energy of polar decoding as a function of gap to capacity. Finally, in Section IX we discuss some open questions and areas of future work.

Notation: We use standard Bachmann-Landau notation in our discussions. The statement $f(x) = O(g(x))$ means that for sufficiently large x, $f(x) \leq cg(x)$ for some positive constant c. The statement $f(x) = \Omega(g(x))$ means that for sufficiently large x, $f(x) \geq cg(x)$ again for some positive constant c. The statement $f(x) = \Theta(g(x))$ means that there are two positive constants b and c such that $b \leq c$ and for sufficiently large x, $bg(x) \leq f(x) \leq cg(x)$. The statement $f(x) = o(g(x))$ means that for sufficiently large x and all $c > 0$, $f(x) < cg(x)$.

We let the symbol $[N] = \{1, 2, \ldots, N\}$ denote the set of integers from 1 to N.

Submitted in part for presentation at the 2016 IEEE International Symposium on Information Theory

Given a set of indices $X, Y \subseteq [N]$, and a vector V of length N, we define the notation $V(X)$ to be the subvector of V formed by the indices in X. As well, given an $N \times N$ matrix A, the notation $A(X, Y)$ refers to the submatrix formed by the rows with indices in X and columns with indices in Y. The notation $A(X)$ refers to the submatrix of A formed by the rows with indices in X and all the columns.

II. Prior Related Work

This paper analyzes the VLSI complexity of polar encoders and decoders. It was recently discovered that the general technique of polar coding was first discovered by Stolte in [9], though these results were never published and the author did not conjecture that this construction reaches capacity. Arikan [7] independently discovered this technique and proved that such codes can reach capacity. Our work in Sections V and VI take inspiration from polar encoding and decoding graphs presented in the Arikan paper.

Our work in Section IV involves a lower bound for circuits that compute polar-encoding functions. The lower bounding technique comes from Thompson [10]. The key lemma needed is Lemma 3, which is analogous to a property of the discrete Fourier transform (DFT) matrix proved by Valiant in [11, Lemma 4] and by Tompa in [12, Lemma 3], though we use a different technique to derive this property.

In Section V we study the butterfly network graph proposed by Arikan [7] for polar decoding. Our key lemma shows that the minimum bisection width of this graph's output nodes is N. This result is similar to the result of [13] which shows that the minimum bisection width of *all* the nodes the butterfly network graph is $2(\sqrt{2}-1)N + o(N) \approx 0.82N$. Because of our circuit lower bounding technique, the minimum bisection width of the output nodes is required, and not all the nodes of the graph, motivating our approach.

In Section VI we show how a mesh network can achieve our encoding and decoding energy lower bounds up to polylogarithmic factors. A mesh network DFT algorithm was proposed by Stevens [14] and shown by Thompson [10] to reach the DFT VLSI complexity lower bounds.

There have been a number of papers on practical VLSI implementations of polar encoders and decoders [15]–[18], though a theoretical analysis of how the energy of such circuits scale has not been performed. However, these results show that practical polar coding circuits compete well with other error control codes, motivating our theoretical analysis.

III. Computational Model

The model we will use in this paper, which we call the Thompson VLSI model, is adapted from Thompson in [10]. The precise model we will use is defined in [3], with a few minor differences. In the previous model, lower bounds were presented in terms of wire width and a technology constant. In this paper, since we are concerned with scaling rules, to avoid unnecessary notation, we will let these constants be 1. As well, in this paper, we let the symbol q represent the switching activity factor (that is, the fraction of the circuit that is active on average per clock cycle). In the previous paper q was implicitly 1. The key circuit parameters we will use are A, the circuit area, and T, the number of clock cycles. Then, the circuit energy is given by $E = qAT$.

Note that in this model circuit nodes can be considered the vertices of a graph, and in a circuit's graph two edges connect a vertex if and only if there is a wire between the two corresponding nodes.

Definition 1. A *bisection* of a set of vertices X of a graph $G = (V, E)$ is a set of edges such that their removal bisects the vertices X; that is, they divide the graph into two disconnected components with vertices V_1 and V_2 such that $||X \cap V_1| - |X \cap V_2|| \leq 1$. The *minimum bisection width* (MBW) of a set of vertices X is the size of the bisection of these vertices that is minimal over all such bisections.

An important lemma relates the minimum bisection width of a graph to its circuit area:

Lemma 1 (Thompson)**.** *All circuits whose corresponding graph has MBW ω have area bounded by:*

$$A \geq \frac{\omega^2}{4}.$$

Proof: See [10]. ∎

Definition 2. A *coding scheme* is a sequence of error control codes of increasing block length N, together with a sequence of *encoding circuits* and *decoding circuits.* It is associated with a particular channel.

Let $P_{\mathrm{e}}(N)$ be the block error probability for the code in the scheme with block length N. Then we can define the following:

Definition 3. [5] An $f(N)$-coding scheme is a coding scheme in which for sufficiently large N, $P_{\mathrm{e}}(N) \leq f(N)$.

Note that this definition classifies coding schemes in terms of their block error probabilities. A "good" coding scheme should thus have this $f(N)$ scale quickly to 0. Using the result of [19] we can see that for every $\epsilon > 0$, there exists an $e^{-N^{1/2-\epsilon}}$-coding schemes using polar codes.

IV. Polar Encoders Lower Bound

In this section we will prove that all polar encoders of rate greater than $1/2$ must have energy that scales as $\Omega\left(N^{3/2}\right)$. The main technical result will be Lemma 3, in which we show a property about the rank of rectangle pairs of the polar encoding generator matrix.

A. Rectangle Pairs

We will consider an $N \times N$ matrix G. We let $R, C \subseteq [N]$.

Definition 4. Let $G(R, C)$ be the submatrix of G formed by selecting the rows with indices in R and the columns with indices in C. We call such an object a *rectangle* of G.

Example 1. Let

$$G = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \\ 9 & 10 & 11 & 12 \\ 13 & 14 & 15 & 16 \end{bmatrix}, R = \{1, 3\}, C = \{2, 4\}$$

then

$$G(R, C) = \begin{bmatrix} 2 & 4 \\ 10 & 12 \end{bmatrix}.$$

Definition 5. If $A \subseteq [N]$, we let the *relative complement* of A be $\bar{A} = [N] \backslash A$, that is, those elements in $\{1, 2, \ldots N\}$ that are not in A. The relevant value of N will depend on context, and when in doubt will be specified clearly.

Again, let $R, C \subseteq \{1, 2, \ldots, N\}$ and G be an $N \times N$ matrix.

Definition 6. A *rectangle pair* is an ordered pair of submatrices of a given matrix G, associated with two sets R and C, defined as $\left(G(R, C), G(\bar{R}, \bar{C})\right)$.

Example 2. If G,R, and C are defined as in Example 1, then the rectangle pair associated with R and C is:

$$\left(\begin{bmatrix} 2 & 4 \\ 10 & 12 \end{bmatrix}, \begin{bmatrix} 5 & 7 \\ 13 & 15 \end{bmatrix}\right)$$

We shall also define:

Definition 7. A *k-row-reduced rectangle pair* of a matrix G is an ordered pair of matrices (X, Y). It is formed by starting with any rectangle pair (A, B) of G and deleting a rows from A to form X and b rows from B to form Y such that $a + b = k$.

Example 3. A 1-row-reduced rectangle pair of the matrix G from Example 2 is

$$\left(\begin{bmatrix} 10 & 12 \end{bmatrix}, \begin{bmatrix} 5 & 7 \\ 13 & 15 \end{bmatrix}\right).$$

which is formed by deleting a row from the first matrix in the rectangle pair

$$\left(\begin{bmatrix} 2 & 4 \\ 10 & 12 \end{bmatrix}, \begin{bmatrix} 5 & 7 \\ 13 & 15 \end{bmatrix}\right)$$

of G.

We will consider the structure of the polar encoding matrix by considering its rectangle pairs.

B. Universal Polar Coding Generator Matrix Properties

Definition 8. The *universal polar coding generator matrix* G_n is a matrix defined recursively by Arikan [7]. Let:

$$F_1 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$$

and

$$F_n = \begin{bmatrix} F_{n-1} & 0 \\ F_{n-1} & F_{n-1} \end{bmatrix}.$$

Then

$$G_n = B_n F_n$$

where B_n is a permutation matrix interpreted as the bit-reversal operator.

The structure of B_n (other than it being a matrix that permutes the rows of F_n) will not matter in the proofs that follow.

We will prove a theorem showing that the sum of the ranks (over the field $\mathbb{F}_2$) of the entries of any rectangle pair of G_N is at least $N/2$. This will imply high VLSI complexity for sufficiently high rate VLSI polar encoders.

We will first consider the ranks of rectangle pairs of F_n. Note that F_n and G_n are $N \times N$ matrices, where $N = 2^n$.

Lemma 2. *Let X be a matrix with entries in a field partitioned as*

$$X = \left[\begin{array}{c|c} A & 0 \\ \hline C & B \end{array}\right],$$

where 0 is a zero submatrix.

Then $\mathrm{rank}(X) \geq \mathrm{rank}(A) + \mathrm{rank}(B)$.

Proof: There are $\mathrm{rank}(A)$ linearly independent rows of A, and $\mathrm{rank}(B)$ linearly independent rows of B. The $\mathrm{rank}(A) + \mathrm{rank}(B)$ rows of X corresponding to these rows are also linearly independent. ∎

Lemma 3. *All rectangle pairs (A_n, B_n) of F_n have* $\mathrm{rank}(A_n) + \mathrm{rank}(B_n) \geq \frac{N}{2}$.

Proof: We will use mathematical induction. Note that

$$F_1 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix},$$

and checking all the possible rectangle pairs (A_1, B_1) of F_1 we have that $\mathrm{rank}(A_1) + \mathrm{rank}(B_1) \geq 1 = \frac{N}{2}$. Now, we assume that for all $k \leq n-1$, for all rectangle pairs of F_k denoted (A_k, B_k):

$$\mathrm{rank}(A_k) + \mathrm{rank}(B_k) \geq \frac{2^k}{2}. \tag{1}$$

Consider a rectangle pair (A_n, B_n) of F_n. Note that A_n can be written as:

$$A_n = \begin{bmatrix} P & 0 \\ Q & S \end{bmatrix}. \tag{2}$$

and B_n can be written as:

$$B_n = \begin{bmatrix} L & 0 \\ M & J \end{bmatrix}.$$

where (P, L) and (S, J) are rectangle pairs of F_{n-1}.

Observe that, by Lemma 2

$$\mathrm{rank}(A_n) \geq \mathrm{rank}(P) + \mathrm{rank}(S) \tag{3}$$

and

$$\mathrm{rank}(B_n) \geq \mathrm{rank}(L) + \mathrm{rank}(J). \tag{4}$$

Since (P, L) and (S, J) are rectangle pairs of F_{n-1}, by the induction hypothesis (1):

$$\mathrm{rank}(P) + \mathrm{rank}(L) \geq \frac{N}{4} \tag{5}$$

and

$$\mathrm{rank}(S) + \mathrm{rank}(J) \geq \frac{N}{4}. \tag{6}$$

Thus, by combining (3) and (4):

$$\mathrm{rank}(A_n) + \mathrm{rank}(B_n) \quad \geq$$
$$\mathrm{rank}(P) + \mathrm{rank}(S) + \mathrm{rank}(L) + \mathrm{rank}(J)$$

By rearranging the right side of this inequality and directly substituting the bounds of (5) and (6) we get:

$$\mathrm{rank}(A_n) + \mathrm{rank}(B_n) \geq \frac{N}{4} + \frac{N}{4} = \frac{N}{2}.$$

∎

Corollary 1. *For any rectangle pair of G_n denoted (A, B):*

$$\mathrm{rank}(A) + \mathrm{rank}(B) \geq \frac{N}{2}.$$

Proof: This Corollary follows by observing that [7] $G_n = B_n F_n$ where B_n is a permutation matrix that permutes the rows of F_n. For every rectangle pair of F_n there is an equivalent rectangle pair of G_n, selected by choosing the same columns and choosing the rows as permuted by B_n. The rectangles forming such a rectangle pair will have the same rows, simply permuted. Thus they have the same row space and thus the same rank. ∎

C. Encoder Circuit Lower Bounds

We consider below a circuit that computes a polar encoding function. Such a function is associated with a set $\mathcal{A}$ of free indices. We denote the vector of free indices $u(\mathcal{A}) \in \{0,1\}^{|\mathcal{A}|}$. It is also associated with a vector of frozen indices $u(\bar{\mathcal{A}}) \in \{0,1\}^{N-|\mathcal{A}|}$.

Definition 9. [7] A *polar encoding function* $f : \{0,1\}^{|\mathcal{A}|} \to \{0,1\}^N$ associated with free indices $\mathcal{A}$ and frozen vector $u(\bar{\mathcal{A}})$ is defined as

$$f(u(\mathcal{A})) = u(\mathcal{A})G_n(\mathcal{A}) + u(\bar{\mathcal{A}})G_n(\bar{\mathcal{A}})$$

where $G_n(\mathcal{A})$ is the submatrix of G_n formed by the rows with indices in $\mathcal{A}$, and addition is performed in $\mathbb{F}_2^N$. Such a function is an encoding function for a code with block length $N = 2^n$ and rate $R = \frac{|\mathcal{A}|}{N}$.

Theorem 1. *The area A and the number of clock cycles T for a circuit that computes a polar encoding function of rate R greater than $\frac{1}{2}$ has area A and number of clock cycles T bounded by*

$$AT^2 \geq \frac{N^2 (1-2R)^2}{64} = \Omega\left(N^2\right) \tag{7}$$

and, if such a circuit has switching activity factor q, its energy is bounded by

$$E \geq q\frac{N^{3/2}(1-2R)}{8} = \Omega\left(N^{3/2}\right). \tag{8}$$

Proof: We will follow a similar line of reasoning used by Thompson in analyzing the complexity of Fourier transform circuits [20]. There are N output bits of the encoder. We label the indices of the output nodes on one side of the bisection L (the left side), and the others R (the right side). The subcircuit containing the output nodes L will have some inputs bits, labelled L_i. Similarly, label the input bits on the right side R_i. We denote the vector of inputs on the left side $u(L_i)$, on the right side $u(R_i)$, and the frozen vector $u(\bar{\mathcal{A}})$. By simply expanding the vector matrix multiplication, we see that the left side of the circuit must compute the vector:

$$\begin{aligned} y(L) &= u(L_i)G_n(L_i, L) + u(R_i)G_n\left(R_i, L\right) \\ &\quad + u(\bar{\mathcal{A}})G_n\left(\bar{\mathcal{A}}, L\right) \end{aligned}$$

and similarly the right side must compute the vector:

$$\begin{aligned} y(R) &= u(L_i)G_n(L_i, R) + u(R_i)G_n\left(R_i, R\right) \\ &\quad + u(\bar{\mathcal{A}})G_n\left(\bar{\mathcal{A}}, R\right). \end{aligned}$$

The subcircuits must compute these values only with the input bits injected into their own input nodes and the bits communicated to them from the other subcircuit (which of course has access to the other input nodes). Note that $\left(G_n\left(R_i, L\right), G_n(L_i, R)\right)$ is an $\left|\bar{\mathcal{A}}\right|$-row-reduced rectangle pair of G_n. Observe from Corollary 1 that the sum of the ranks of these matrices must be at least $\frac{N}{2} - \left|\bar{\mathcal{A}}\right| = \frac{N}{2} - (1-R)N$, which is greater than 0 because $R > 1/2$. Thus, at least $\frac{N}{2} - (1-R)N$ bits must be communicated across this bisection during the computation. If the circuit has MBW of output bits ω, since at each clock cycle at most 2ω bits can be communicated across the bisection:

$$T \geq \frac{\frac{N}{2} - (1-R)N}{2\omega}. \tag{9}$$

By Lemma 1, we have $A \geq \frac{\omega^2}{4}$ and thus, combining this with (9) implies

$$AT^2 \geq \frac{N^2 (2R-1)^2}{64}. \tag{10}$$

Also note that

$$A \geq N$$

and thus, combining this inequality with (10) and taking the square root we get

$$E = qAT \geq q\frac{N^{1.5}(2R-1)}{8}.$$

■

V. Arikan Successive Cancellation Polar Decoding Scheme

In Arikan's original paper on polar coding [7], the author presented a Turing-time complexity $O(N \log N)$ algorithm for computing successive cancellation decoding of polar codes. In this section, we provide a definition of a polar decoder based on Arikan's [7] paper, and show that such circuits, when implemented with output nodes arranged in a rectangular grid, take energy at least $\Omega(N^{3/2})$.

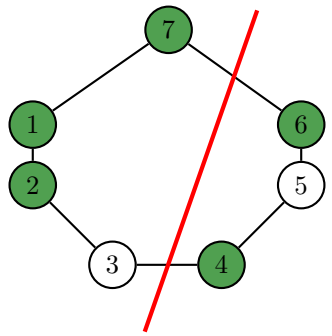

Figure 1: Example of a graph with a 2-partition of the shaded nodes. The labelled partitioning line is also a 1-partition of the white nodes. Inspection shows that such m-partitions are minimal. Therefore, the 2-section width of the shaded nodes in this graph is 2, as is the 1-section width of the white nodes.

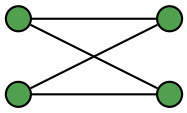

Figure 2: The base case: the decoding graph P_1.

A. Decoding Complexity Lower Bound

Below we consider a generalization of the minimum bisection width of a set of vertices, where instead of dividing the set of vertices into two sets of equal size, instead we divide the vertices of a set into two sets, where the size of one of these sets is fixed.

Definition 10. Given a graph $G = (V, E)$, an *m-partition* of a set of vertices $X \subseteq V$ is an ordered pair $(A, V \backslash A)$ in which $A \subseteq V$ and $|A \cap X| = m$. The *width* of this partition is the size of the set of edges connecting the vertices in A with $V \backslash A$. The *m-section width* of a set X of vertices of a graph G is the minimum width over all the graph's m-partitions of X.

In Figure 1 we give an simple example of a 2-partition of a subset of edges of a graph.

Note that if a graph $G = (V, E)$ has $2m$ vertices, then the m-section width of V is the same as the graph's minimum bisection width.

Definition 11. An *n-polar decoding graph,* denoted P_n, is a generalization of the graph presented by Arikan in [7] describing the communication graph of a polar decoding algorithm. It is defined recursively in Figures 2 and 3. For the base case, the 1-polar decoding graph is the bowtie shaped graph given in Figure 2. An n-polar decoding graph consists of 2^n nodes called *symbol nodes* connected to two copies of the $(n-1)$-polar decoding graphs as shown in Figure 3.

We shall study the structure of an m-partition of a polar decoding graph.

We call the nodes in the left-most column in the graph of Figure 3 *symbol nodes*. Note that for any n-polar decoding graph, there exist 2^n symbol nodes, as well as symbol nodes of the two $(n-1)$-polar decoding subgraphs that form the graph. Let the set of symbol nodes of the larger graph be labelled A, the symbol nodes of one of the subgraphs be B, and the other symbol nodes C. This labelling of sets is visible in Figure 3.

Note by inspection that that the bipartite subgraph connecting the nodes of A with $B \cup C$ consists of *bowties*, subgraphs containing two vertices from A and a vertex from B and C. An example bowtie is labelled in Figure 3. A bowtie object is a subgraph of the polar decoding graph associated with a particular partition. We classify these bowties according to how the particular partitioning line divides the nodes in the graph.

We now consider a minimum m-partition of A. Such a partition divides the set of vertices into two subsets, which we will call the top half and the bottom half. We can consider how the 2^{n-1} bowties connecting A with B and C are split by the m-partition. We divide such bowties into three categories: split bowties, contained bowties, and crossing bowties.

Definition 12. A *split bowtie* is a bowtie in which one element of A is in the upper half, and one in the lower half. Two examples are given in Fig. 4a.

Note that a split bowtie has 2 edges crossing the m-partition.

Definition 13. A *contained bowtie* is one in which all vertices are either in the top or bottom half. An example is given in Figure 4b.

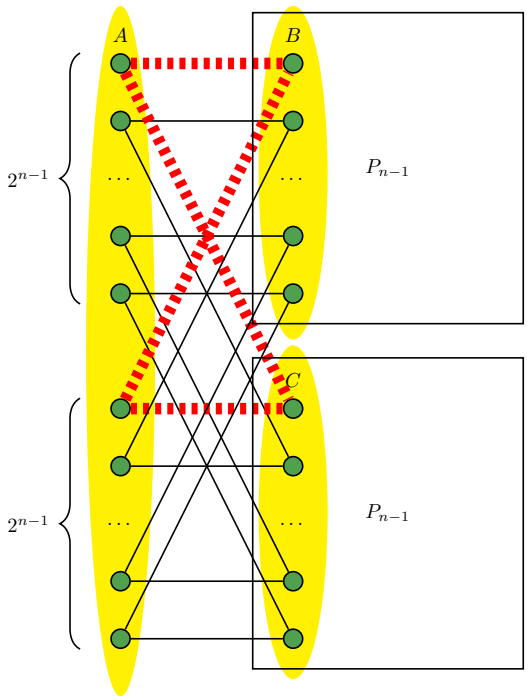

Figure 3: A diagram of the recursive structure of a polar decoding graph P_n. The vertices on the left (outlined in a shaded ellipse and labelled A) indicate the symbol nodes of the graph. The vertices on the right in the two boxes indicate the symbol nodes of the two subgraphs (outlined and labelled B and C) which are the smaller polar decoding graphs P_{n-1}. An example of a subgraph that is a bowtie is indicated with edges drawn with a thick, dashed line.

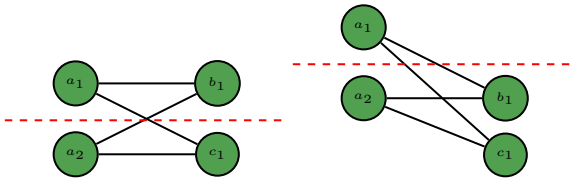

(a) Two examples of split bowties. Split bowties are bowties in which the two nodes in A are on opposite sides of the partition line. It does not matter where the other nodes lie.

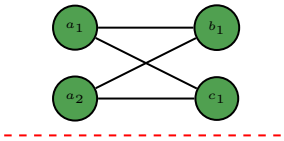

(b) An example of a contained bowtie. Such a bowtie occurs when all the nodes of a bowtie are on the same side of the partition.

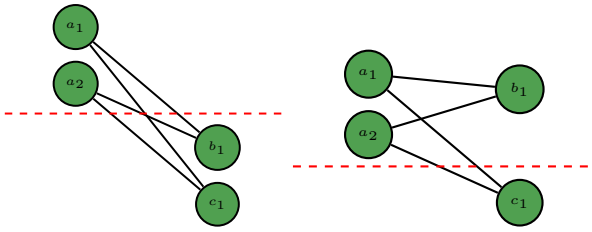

(c) Two examples of crossing bowties. In such bowties the two nodes in A are on the same side of the partition, and at least one of the other nodes in the bowtie is on the other side.

Figure 4: Diagrams of split, contained, and crossing bowties. Note that a bowtie is an object associated with a polar decoding graph and a particular partition. In each figure, the nodes labelled a_1 and a_2 are nodes from the set A. Similarly, the nodes labelled b_1 are nodes from the set B and those labelled c_1 are from the set C. The dashed line indicates the relative position of the partitioning line.

Note that a contained bowtie has no edges crossing the m-secting cut.

Definition 14. A *crossing bowtie has two nodes of A on one side of the partition line, and at least one of the nodes in B or C on the other side of the partition line.*

It should be clear that for any partition of the polar decoding graph, all bowties are either split, contained, or crossing bowties.

We let G be an arbitrary polar decoding graph.

Lemma 4. *Consider an n-polar decoder graph, and let $0 \leq m \leq 2^{n-1}$. Such a graph has a minimum m-section partition of the symbol nodes of G which contains only split bowties and contained bowties.*

Proof: We provide an exchange argument. For a particular minimal m-partition, we argue that nodes in a crossing bowtie can be moved to another side, resulting in a contained bowtie without increasing the number of edges crossing the partition line. By examining the left side of Figure 4c, there are 4 edges crossing the partition line. For crossing bowties of this type, moving vertices b_1 and c_1 to the side containing the vertices in A decreases the edges of this bowtie crossing the partitioning line by 4. The degree of any vertex in this graph is at most 4, thus this can, at worst, result in 4 new edges crossing the partition (which would result from the two extra edges each on vertices b_1 and c_1 now crossing the partition line). A similar argument can be made for crossing bowties like those in the right side of Figure 4c. Thus, any minimum m-section partition of the output nodes containing crossing bowties can be modified to one that contains no crossing bowties. ■

We shall prove a lemma regarding the m-section width of an N-polar decoding communication graph.

Lemma 5. *Let $0 \leq m \leq 2^{n-1}$. Then the m-section width of the symbol nodes of an n-polar decoding graph is at least $2m$.*

Proof: We shall prove this by induction. For the base case, the $n = 1$ polar decoding graph given in Figure 2 can be shown by inspection to satisfy the lemma by simply checking the width of all 0 and 1-partitions.

We shall assume that the m-section width, where $0 \leq m \leq 2^{k-2}$, of a $(k-1)$-polar decoding graph is $2m$.

We consider a minimal m-partition that contains only split and contained bowties that exists by Lemma 4. Denote one half of the partition the upper half and the other the lower half. Without loss of generality we assume the upper half contains m nodes of A. Let the number of contained bowties in the upper half be C_{upper}, and the number in the lower half C_{lower}. Let the number of split bowties be S. Since the upper half contains m nodes of A, then

$$2C_{\text{upper}} + S = m. \tag{11}$$

Note that the number of contained bowties on the lower half must at least equal the number on the upper half, since the number of symbol nodes on the lower half must equal or exceed the number in the upper half. Thus, there must be at least m elements of B and m elements of C on both side of the partition. As well, at least one of these sides cannot have more than $N/4$ elements (since each of these sets contains only $N/2$ elements in total). Thus, there is an x-partition of both B and C induced by the partition, where $C_{\text{upper}} \leq x \leq N/4$. Thus, each of these partitions must have at least $2C_{\text{upper}}$ edges crossing the partition line by the induction hypothesis. In addition, there are 2 edges crossing the partition line for each split bowtie (easily observed by inspecting Figure 4a). Thus, the number of edges crossing the partition line is at least

$$\text{Edges crossing} \geq 4C_{\text{upper}} + 2S = 2m$$

where we have applied (11), proving the theorem. ■

We will consider algorithms whose communication graph is based on the polar decoding graph. However, bits corresponding to certain symbol nodes which are frozen obviously do not need to have their own node in a communication graph. Thus we consider a frozen-bit polar decoding graph.

Definition 15. An *n-frozen bit polar decoding graph* associated with frozen bit indices $\bar{\mathcal{A}}$ is a graph obtained by deleting the symbol nodes corresponding to $\bar{\mathcal{A}}$ from P_n and also the edges to which they are connected. The symbol nodes that remain are called the *unfrozen nodes*. Such a graph is a decoding graph for a rate $R = 1 - \frac{|\bar{\mathcal{A}}|}{N}$ code.

Note that this is a natural simplification of the polar decoding graph once frozen bits are considered. However, this is not the only possible simplification. In this paper we only consider simplifications that involve deletion of the nodes corresponding to the frozen bits.

Once the symbol nodes corresponding to frozen bits are deleted, we then consider the bisection width of the remaining symbol nodes.

Corollary 2. *The minimum bisection width ω of the unfrozen nodes of any n-frozen bit polar decoding graph in which $R > 2/3$ is at least:*

$$\omega \geq N(3R - 2)$$

Proof: Suppose not. Then, consider the unfrozen symbol nodes minimal bisection with $W < N(3R-2)$ nodes crossing it. Now, add the frozen symbol nodes and their edges back to this graph. There are at most $(1-R)N$ such nodes to be added,

and thus they can increase the number of edges in the graph by at most $2(1-R)N$. At most all these edges can cross the partition line, and thus this partition line can have at most $W + 2(1-R) < RN$ edges crossing it. Note that the partition line forms an m-partition of all the symbol nodes, where $m \geq RN/2$, But, by Lemma 5, any such m-partition must have at least RN edges crossing it, resulting in a contradiction. ■

B. Decoder VLSI Lower Bounds

Note that a Thompson circuit is associated with a graph. In the course of a computation, messages will be passed from node to node in the circuit. Each binary message passed corresponds to another edge in the computation's *communication graph*. We will define a polar decoder (a type of circuit) in terms of a circuit's communication graph.

We first define contracting vertices and subdividing edges.

Definition 16. Any two vertices of a graph u and v can be *contracted* by joining two vertices together, resulting in a new vertex v'. The vertex v' will have all the edges corresponding to both u and v with the exception of any edges that joined u and v in the original graph, which are deleted.

Definition 17. An edge of a graph can be *subdivided* by being replaced by two edges with a vertex in the middle.

Consider a frozen polar decoding graph. Such a graph has RN symbol nodes. Now, consider a process in which vertices and edges are added to this graph, and some edges are subdivided and some vertices are contracted, but no two symbol nodes are contracted.

Definition 18. Such a graph is called a *polar decoding communication graph*. Obviously, such a graph has RN symbol nodes, which are the nodes that correspond to the symbol nodes of the original frozen polar decoding graph. We call these nodes the graph's *output nodes.*

Lemma 6. *All polar decoding communication graphs obtained from a rate $R > 2/3$ frozen polar decoding graph have minimum bisection width of output nodes at least $N(3R-2)$.*

Proof: First, observe from Corollary 2 that the minimum bisection width of the symbol nodes of the original frozen polar decoding graph is at least $N(3R-2)$. Suppose that the polar decoding communication graph obtained from this graph has minimum bisection width less than this. Then consider the partition with $W < N(3R-2)$ edges crossing the partition line. We can reverse the operations that obtained the the communication graph and this can only decrease the number of edges crossing the partition line, resulting in a bisection of less than $N(3R-2)$ for the original frozen polar decoding graph, a contradiction. To see this, consider what happens when a vertex v' that was contracted has this operation reversed, resulting in two vertices with u and v with (possibly) an edge between them. Place u and v on the same side of the partition previously occupied by v'. It is then obvious that any edge between them cannot increase the width of this partition. It is also obvious that deleting any vertices and edges that were added can only decrease the the number of edges crossing the partition line. The same is true for reversing a subdividing process. ■

Definition 19. An *Arikan polar decoding circuit* is a circuit whose associated communicaiton graph is also a polar decoding communication graph.

Note that in our model, a Thompson circuit is created by placing nodes and edges each in a grid of squares of side length 1. Consider placing a grid of (possibly larger) squares with integer side length on top of any Thompson circuit and with sides aligned to the smaller grid of squares defining the VLSI circuit. Then each square in the grid will contain some output nodes.

Definition 20. A *rectangle grid output circuit* with N output nodes is a circuit in which there is a grid of squares that can be placed upon the circuit, and there is a $\sqrt{N} \times \sqrt{N}$ array of larger grid squares, each which contains exactly one output node.

Example 4. The mesh network defined in Section VI is an example of a rectangle grid output circuit.

We suspect that our scaling rule lower bounds for polar decoders would extend to implementations that are not necessarily rectangle grid output circuits, however for simplicity we only present our results for such circuits. A generalization of the following lemma to a broader class of circuits would be sufficient:

Lemma 7. *All rectangle grid output circuits with $\Theta(N)$ output nodes and with a communication graph with minimum bisection width ω have energy $E \geq \Omega\left(\omega\sqrt{N}\right)$.*

Proof: See Appendix A. ■

Note that the above lemma does not assume a constant switching activity factor.

Theorem 2. *All rectangle grid output, Arikan polar decoding circuits have energy that scales at least as $\Omega\left(N^{3/2}\right)$.*

Proof: This flows directly from Lemma 6 and Lemma 7. ■

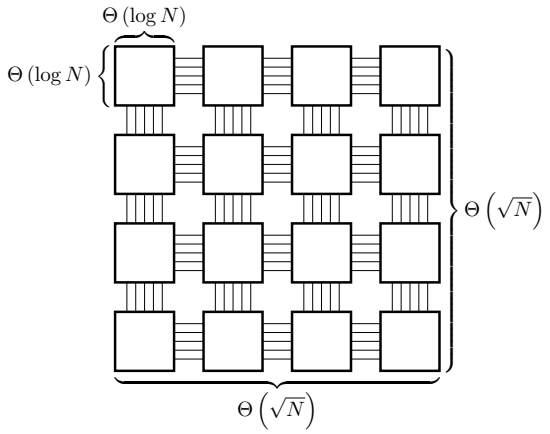

Figure 5: Diagram of a mesh network. A mesh network consists of a grid of $\sqrt{N} \times \sqrt{N}$ processor nodes, each with area that scales as $\Theta(\log^2 N)$. Each node is connected to its at most 4 neighbors with $\Theta(\log N)$ wires.

It is now trivial to observe that this lower bound can be reached up to a polylogarithmic factor on a mesh network which we discuss in the Section VI-D.

VI. Upper Bounds

A. Mesh Network

We will show that, up to polylogarithmic factors, the lower bounds on the energy of polar encoding and decoding complexity can be reached. The mesh network topology that we present to meet these bounds derives from Thompson [10].

A mesh network consists of a grid of processor nodes. Each processor node is capable of sending and receiving messages to its adjacent nodes. As shown in Figure 5, each node has area that scales as $\Theta(\log^2 N)$, and consists of a processor and memory. The processor takes area that remains constant with increasing circuit size, so the amount of memory in each node can scale as $\Theta(\log^2 N)$. Each processor node must also contain instructions on what each node is to compute. The length of the instructions obviously cannot be longer than $O(\log^2 N)$. As is clear from the diagram, each processor node is connected to up to 4 other processor nodes with $\Theta(\log N)$ wires.

A computation on a mesh network consists of a message-passing procedure and the computation that the nodes are to perform on the messages they receive. A typical message consists of an *address*, an encoding of the node to which the message is to be sent, and its *content*, the information meant to be sent. Since there are $\Theta(N)$ processor nodes, an addressing scheme with $\Theta(\log N)$ bits per address is sufficient. As well, a constant-sized message and the address of the node to which the message is to be sent can be passed between adjacent nodes in a single clock cycle, because the width of the wires connecting such nodes also scales as $\Theta(\log N)$. (Actually, messages with size that scales logarithmically in N can be sent).

Multiple messages may be sent simultaneously in a mesh network, but in order for an algorithm to be *valid* for a mesh-network, it must be that no computational node is required to store more than $O(\log^2 N)$ bits in its memory. As well, for an algorithm to be valid, it must also avoid message-passing *conflicts:* two messages cannot be passed to the same node at the same time. Thus, given a mesh-network algorithm, we must show that its message passing order does not result in any conflicts.

In the section below we provide a message-passing procedure for computing polar encoding.

B. Encoding

Arikan provides a method for computing polar encoding that naturally lends itself to implementation on a mesh network. See Figure 3 for a graphical representation of the Arikan method. In the Arikan method, the input nodes are on the left side of the graph. As well, for polar encoding, some of these nodes represent frozen bits. In the encoding algorithm, messages move left to right. The input nodes (and frozen bit nodes) first pass their bits to the node to which they are connected in the adjacent column of nodes. The nodes in this column proceed to compute the modulo 2 sum of their inputs, and pass the result to their adjacent nodes on the right. This continues until the final column is reached, resulting in the codeword.

In our proposed mesh-network implementation, each of the N processor nodes corresponds to a row of nodes in the encoding graph of Figure 6. Obviously, if each message (which corresponds to an edge in the graph) is to be sent one-by one, such an order of the message-passing procedure would avoid conflicts and would be easily be implementable on a mesh network. (In fact, this is the way we propose to do decoding). However, much of the computation for encoding can be done in parallel. We show in Appendix B how a constant fraction of the messages corresponding to edges connecting nodes in adjacent columns can be sent simultaneously in a way that avoids conflicts. We suspect that most sufficiently sparse communication graphs can

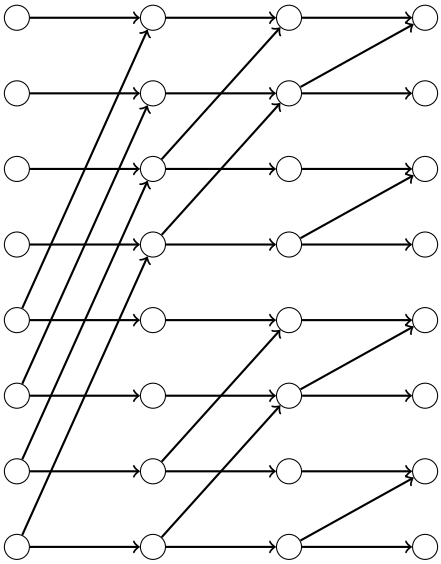

Figure 6: Example of an encoding graph for $N = 2^3 = 8$ taken directly from Arikan [7]. The leftmost column of nodes are the input bits, sending their bits to to their adjacent nodes. Upon receiving these bits, the nodes in the second column compute the $\mod 2$ sum of their inputs, and then pass this result to their adjacent nodes. This procedure naturally suits implementation on a mesh network. Such encoding reaches the energy complexity lower bounds in polar encoding up to polylogarithmic factors.

be implemented on a mesh network in a way that avoids conflicts, however a general analysis of this is beyond the scope of this paper.

C. Analysis of Mesh Network Encoding Algorithm Complexity

Note that there at $\Theta(\log_2 N)$ stages. Suppose the number of clock cycles used by an individual node for reading an address and computing which direction to send its message is t_R. Suppose that the complexity for computing parity of the received bit with the current bit is t_P. At each stage, the number of hops between processor nodes is at most $O(\sqrt{N})$. There are $\Theta(\log N)$ stages. Thus, the number of clock cycles required is

$$T = \Theta\left(\log N\left(\sqrt{N}T_{\mathrm{R}} + T_{\mathrm{P}}\right)\right)$$

The computation of parity can obviously be done in time $\Theta(1)$. The routing requires computing which direction to "send" the message: up, right, or left. This can easily be accomplished in $\Theta\left(\log N\right)$ time (that is, proportional to the length of the address).

The proposed algorithm also uses roughly the same fraction of node each clock cycle, so we can assume for scaling rules the switching activity factor (q) is constant. The area of such a circuit scales as $A = \Theta(\log^2 N)$. Thus:

$$E = qAT = \Theta(N^{3/2}\log^4 N)$$

D. Decoding Mesh Network

Clearly, because of the requirement of successively computing each estimate in polar decoding, asymptotically the number of clock cycles for a polar decoding scheme must scale at least as $\Omega(N)$. A fully parallel polar decoder thus must have area-time complexity at least $\Omega(N^2)$.

However, the algorithm proposed by Arikan [7] that takes time complexity $O(N\log N)$ can easily be implemented on a mesh network. Each node of the mesh network corresponds to a row of nodes of the graph in Figure 3. As described by Arikan [7], a depth first message-passing procedure between the nodes of the graph can compute the polar decoding in Turing time complexity $O(N\log N)$. The distance between any two nodes in a mesh network is not greater than $O(N^{1/2})$. Thus, decoding on a mesh network takes $A = \Theta\left(N\log^2 N\right)$ and $T = \Theta(N^{3/2}\log^2 N)$, where the algorithm takes $O(N\log N)$ steps, and $O(N^{1/2}\log N)$ time to do the message passing. Since a fully parallel decoding algorithm requires only a single processing node to be active at a given time, the switching activity factor of this scheme scales as $\Theta\left(1/N\right)$. Thus the energy of the computation scales as $E = O\left(N^{3/2}\log^4 N\right)$.

VII. Generalized Polar Coding on a Mesh Network

Arikan [7] proposes a generalization of polar codes in which the generator matrix is no longer G_1 as defined in Section IV. Hassani *et al.* [8] analyze such schemes and show that there exist generating matrices in which for sufficiently large N, $P_{\mathrm{N}} \leq e^{-n^{1-\epsilon}}$ for any $\epsilon > 0$. That is, they are $e^{-n^{1-\epsilon}}$-coding schemes. By [1, Theorem 1], such circuits must have bit-meters energy that scales as $E \geq \Omega(N^{3/2-\epsilon/2})$. Both [7] and [8] argue that such schemes have $O(N \log N)$ time complexity algorithms for decoding. When implemented on a mesh network, such algorithms would take energy complexity $\Theta(N^{3/2} \log^4 N)$ for the same reasons described in Section VI-D. Thus, we can say that the general lower bounds can be almost reached for such close-to-exponential coding schemes (that is, within a factor of $N^{\epsilon}\mathrm{polylog}(N)$). Such a scheme would have a switching activity factor q that scales as $\Theta(1/N)$. Such a circuit would take number of clock cycles that scales as $T = \Theta(N^{3/2} \log^2 N)$. Note however that the Thompson complexity analysis of [5] suggests a lower bound $T(N)$ of $\Omega(N^{1/2})$, and thus this method does not simultaneously reach these energy and time lower bounds.

We conjecture that the Thompson model lower bounds for time and energy of [5] for encoding can be reached with polar encoding up to an $N^{\epsilon}\mathrm{polylog}(N)$ factor for small ϵ. This is because parallelization of the encoding procedure is possible. However, such a parallel algorithm must avoid conflicts on the mesh network, so proving this remains an open question.

VIII. Energy Scaling as Function of gap to Capacity

In this section, we consider how the energy of polar codes scales as capacity is approached.

Definition 21. For a particular code, let $\chi = \frac{1}{1-\frac{R}{C}}$ be the *reciprocal gap to capacity*.

Note that as rate approaches capacity, χ approaches infinity.

Guruswami *et al.* [21] show that as a function of reciprocal gap to capacity, the block length required to achieve a set probability of block error P_{e} for polar codes scale as $N = O\left(\chi^{\mu}\right)$ for some value μ; that is, the block length scales polynomially in the reciprocal gap to capacity. Hassani *et al.* [8] show that $3.55 \leq \mu \leq 6$. Goldin *et al.* [22] improve the upper bound on μ to 5.7 and Mondelli *et al.* [23] further improve the upper bound to 4.7. These bounds, combined with Theorems 1 and 2 and the discussion in Section VI that bound energy of encoding and decoding in terms of N by:

$$\Omega\left(N^{1.5}\right) \leq E_{\mathrm{comp}} \leq O\left(N^{1.5} \log^4 N\right)$$

imply an obvious corollary.

Corollary 3. *The energy for polar encoders with reciprocal gap to capacity χ and a set probability of block error, in which $C > \frac{1}{2}$, is bounded by:*

$$\Omega\left(\chi^{5.325}\right) \leq E_{\mathrm{comp}} \leq O\left(\chi^{7.05} \log^4\left(\chi\right)\right) \tag{12}$$

with decoding energy bounded similarly.

Note that polar codes are $e^{-N^{\frac{1}{2}-\epsilon}}$-codes [24] for any $\epsilon > 0$. We can apply the well known general lower bound on block length of any code as a function of fraction of capacity [25] of $N \geq \Omega(\chi^2)$ and the general lower bound of [5, Theorem 1] to show that all $e^{-N^{\frac{1}{2}-\epsilon}}$-encoding and decoding schemes have energy bounded by:

$$E \geq \Omega\left(N^{\frac{5}{4}-\epsilon}\right) \geq \chi^{2.5}$$

When contrasted with the lower bounds of (12), this illuminates a gap between general lower bounds and that which is achievable through polar coding.

IX. Future Work

Extending the lower bound results to polar codes with rates less than $1/2$ for encoder lower bounds and less than $2/3$ for decoding lower bounds is an obvious area of future work. Moreover, we have not actually implemented the mesh network topology analyzed in Section VI, so this also remains an area of future work.

Our decoding lower bounds are for algorithms based on graphs obtained from the butterfly network decoding graph suggested by Arikan [7]. Our lower bound does not necessarily include all decoders based on successive cancellation decoding for polar codes. A particular challenge for generalizing this result is in defining precisely what a circuit that performs a polar decoding algorithm actually is, and thus this remains an area of future work.

Currently, our results apply to polar codes as defined in [7]. Arikan also suggested that other matrices can be used to generate polar codes with similar decoding and encoding complexity. Korada *et al.* [26] showed that there exist matrices which generate polar codes that have block error probability close to $O(2^{-N})$. Our work in [5] shows that such codes necessarily have encoding and decoding energy that scales as $\Omega\left(n^{1.5}\right)$, and implementation over a mesh network as described in Section VI allows for this encoding and decoding in a way that scales as $O(n^{1.5}\mathrm{polylog}\left(n\right))$. Thus, this means polar codes are energy-optimal over all exponentially low probability of error decoders.

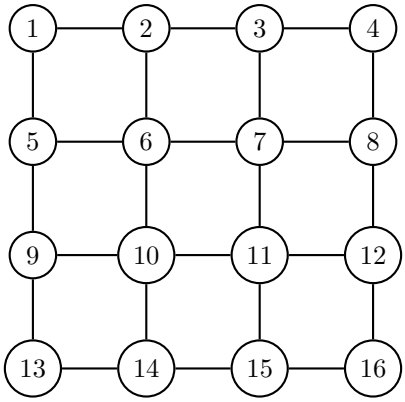

Figure 7: The "raster-scan ordering" proposed for the nodes $i = 1$ to 16 for an $N = 16$ polar encoding circuit. Note that for general N these may not fit perfectly on a square. In that case, we propose using dimensions that are as close to square-like as possible. In Appendix B we show that labelling nodes like this and applying our proposed message-passing procedure for polar encoding avoids conflicts and thus is valid.

However, a full understanding of VLSI complexity of polar encoding should relate the generator matrix to their asymptotic VLSI complexity, including energy, time, and area tradeoffs. Our lower bound approach here works because of the structure of the G_1 matrix, allowing us to derive Lemma 2. This lemma is used in a key step in the inductive proof showing that polar encoding has high VLSI complexity. However, this approach does not seem to extend to different generator matrices; thus we cannot rule out that there are some polarizing matrices with encoding energy that scales better than $\Omega(n^{1.5})$. These may possibly come at the cost of higher error probability. For engineering purposes, however, this may be a worthwhile tradeoff, and the analysis of polar coding may be a tool to characterize achievable coding schemes that reach different points in this fundamental tradeoff.

X. Appendix 1

Appendix A

Proof: (Of Lemma 7) Note that across any bisection of the communication graph of a given circuit, there are at least ω edges that cross that bisection, with total length at least ω. We can construct a single bisection with a line through the middle of the nodes. Then we can construct another bisection by shifting this line left one unit, and sweeping in a parallel line from the right. Such a bisection has half the nodes in between the two lines and half the nodes outside the two lines. We can then shift these two lines left one unit again. We can do this on the order of $\sqrt{N}$ times and each time the edges crossing the bisecting lines must be at least ω, and each time the bisecting lines are in a different position, thus for each bisection the parts of the lines that cross the bisecting lines are different. In total there must be at least $\Omega\left(\sqrt{N}\omega\right)$ total distance of the lines and thus this amount of total energy. ■

Appendix B
Mesh Network Encoding Procedure

We implement the message-passing procedure on a mesh network using $n = \log N$ stages, one stage for each pair of adjacent columns in the encoding graph of Figure 6. Label the input nodes of the encoding graph in order, starting from the top, from 1 to N. For each input node i of the encoding graph, associate a mesh network processor node i. This processor node is to perform all the computations and message passings of the graph nodes in the same row as its associated input node. Place the processor nodes on the $\sqrt{N} \times \sqrt{N}$ mesh network in the order shown in Figure. 7. We call such an ordering a *raster scan ordering*. Note that by inspecting Figure 6, in the jth stage of the procedure, each processor node i must pass messages to node $i - 2^{n-j}$.

There is some ambiguity in our definition of the operation of the mesh network: when a message is received by a processor node, in which direction should the message be sent? It is clear that if the message is located "above and to the left" of the node, the message should be passed either to the left or above. We shall adopt the convention that a node shall choose the relevant up or down direction before deciding to send the node left or right, which occurs at what we will call the *target row* (that is, the row containing the computational node to which the message is sent).

We can label each row of the mesh network, and thus some nodes will be on even rows, and some nodes will be on odd rows. In our proposed procedure, each stage of the encoding will be divided into two message-passing steps: the "even-row" passing step and the "odd-row" passing step. More precisely, at the jth stage of the encoding procedure, only nodes i on even

rows that are to send their bits to node $i - 2^{n-j-1}$ shall do so. Then, the appropriate nodes on odd rows are to do the same. We claim this procedure avoid conflicts (that is, no two messages will be sent to the same node simultaneously).

Definition 22. A *constant send-back procedure* is a message passing procedure defined on a mesh network with nodes labelled according to the raster-scan ordering in which a set of nodes, indexed i, each simultaneously send a message to node $i - m$, for some $m > 0$.

The "even row" sending step and the "odd row" sending step of the procedure we propose for polar encoding is obviously a special case of this procedure.

Lemma 8. *In any constant send-back procedure, conflicts can only occur with messages originating on different rows.*

Proof: Consider two messages originating on the same row. Since our convention is to send nodes "up" until deciding to send them left or right, a conflict between these two messages can only occur on the target row in which one message is sent left, and the other sent right. However, because of the ordering of the processor nodes, and the fact that we are considering a constant send-back procedure, upon reaching the target row, these messages must be sent in the same direction, otherwise they are not addressed to nodes a constant value less than their index. If these nodes do not have the same target rows, then the lemma flows trivially. ■

Lemma 9. *Messages originating on rows spaced at least* 2 *rows apart cannot have the same target rows and can not conflict.*

Proof: Let x be the number of processor nodes in each row. Clearly, the spacing between two nodes at least 2 rows apart is at least $x + 1$ (occurring when the lesser indexed node is at the far right, and the greater indexed node is on the far left). Suppose their target node was on the same row. The spacing between these two target indices must at least be $x + 1$, but there are only x elements on each row. Those, messages on rows spaced two or more apart cannot conflict in their target rows. It may be possible for them to conflict where one message has reached a target row, and is travelling left or right, and another is travelling up. However, the message travelling up must have originated on a row below the left or right-going node. In a constant send-back procedure such a message cannot have a target row at the same level or higher than the other node, so this cannot occur. ■

Since in each step, at each stage of our proposed procedure, no two simultaneously sent messages originating on adjacent rows, combining Lemmas 8 and 9 confirms that our proposed procedure has no conflicts.

This particular message passing ordering is done entirely to prove that the area-time complexity *scaling* of this mesh scheme is close to the lower bound. It is likely that in any practical implementation a more efficient message passing procedure exists (though it will likely be more efficient only up to some constant factor).

REFERENCES

[1] P. Grover, "Information friction and its implications on minimum energy required for communication," *IEEE Trans. Inf. Theory*, vol. 61, no. 2, pp. 895–907, Feb 2015.

[2] P. Grover, A. Goldsmith, and A. Sahai, "Fundamental limits on the power consumption of encoding and decoding," in *Proc. 2012 IEEE Int. Symp. Info. Theory*, 2012, pp. 2716–2720.

[3] C. G. Blake and F. R. Kschischang, "Energy consumption of VLSI decoders," *IEEE Trans. Inf. Theory*, vol. 61, no. 6, pp. 3185–3198, June 2015.

[4] ——, "On the energy complexity of LDPC decoder circuits," *CoRR*, vol. abs/1502.07999, Feb. 2015. [Online]. Available: http://arxiv.org/abs/1502.07999

[5] ——, "Energy, latency, and reliability tradeoffs in coding circuits," 2016, in preparation.

[6] K. Ganesan, P. Grover, J. Rabaey, and A. Goldsmith, "Towards approaching total-power-capacity: Transmit and decoding power minimization for LDPC codes," *CoRR*, vol. abs/1504.01019, April 2015. [Online]. Available: http://arxiv.org/abs/1504.01019

[7] E. Arikan, "Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels," *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3051–3073, July 2009.

[8] S. Hassani, K. Alishahi, and R. Urbanke, "Finite-length scaling for polar codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 10, pp. 5875–5898, July 2014.

[9] N. Stolte, "Rekursive codes mit der plotkin-konstruktion und ihre decodierung," Ph.D. dissertation, TU Darmstadt, 2002.

[10] C. D. Thompson, "A complexity theory for VLSI," Ph.D. Thesis, Carnegie-Mellon, 1980.

[11] L. G. Valiant, "Graph-theoretic properties in computational complexity," *Journal of Computation and Systems Sciences*, vol. 13, pp. 278–285, 1976.

[12] M. Tompa, "Time-space tradeoffs for computing functions, using connectivity properties of their circuits," *Journal of Computer and System Sciences*, vol. 20, no. 2, pp. 118–132, April 1980.

[13] C. Bornstein, A. Litman, B. Maggs, R. Sitaraman, and T. Yatzkar, "On the bisection width and expansion of butterfly networks," in *Parallel Processing Symposium, 1998. IPPS/SPDP 1998. Proceedings of the First Merged International ... and Symposium on Parallel and Distributed Processing 1998*, Mar 1998, pp. 144–150.

[14] J. E. Stevens, *A fast fourier transform subroutine for ILLIAC IV*. Urbana, IL, USA: Center for Advanced Computation, University of Illinois at Urbana-Champaign, 1971.

[15] O. Dizdar and E. Arikan, "A high-throughput energy-efficient implementation of successive cancellation decoder for polar codes using combinational logic," *CoRR*, vol. abs/1412.3829, 2015. [Online]. Available: http://arxiv.org/abs/1412.3829

[16] A. Balatsoukas-Stimming, A. Raymond, W. Gross, and A. Burg, "Hardware architecture for list successive cancellation decoding of polar codes," *IEEE Trans. Circuits and Systems II: Express Briefs*, vol. 61, no. 8, pp. 609–613, Aug 2014.

[17] A. Alamdar-Yazdi and F. R. Kschischang, "A simplified successive-cancellation decoder for polar codes," *Communications Letters, IEEE*, vol. 15, no. 12, pp. 1378–1380, December 2011.

[18] C. Leroux, I. Tal, A. Vardy, and W. Gross, "Hardware architectures for successive cancellation decoding of polar codes," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, May 2011, pp. 1665–1668.

[19] E. Arikan and I. Telatar, "On the rate of channel polarization," in *IEEE Int. Symp. Inf. Theory*, June 2009, pp. 1493–1495.

[20] C. D. Thompson, "Area-time complexity for VLSI," *Proc. 11th Ann. ACM Symp. Theory of Comput.*, pp. 81–88, 1979.

[21] V. Guruswami and P. Xia, "Polar codes: Speed of polarization and polynomial gap to capacity," *IEEE Trans. Inf. Theory*, vol. 61, no. 1, pp. 3–16, Jan 2015.
[22] D. Goldin and D. Burshtein, "Improved bounds on the finite-length scaling of polar codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 6966–6978, Nov 2014.
[23] M. Mondelli, S. Hassani, and R. Urbanke, "Unified scaling of polar codes: Error exponent, scaling exponent, moderate deviations, and error floors," *CoRR*, vol. abs/1501.02444, 2015. [Online]. Available: http://arxiv.org/abs/1501.02444
[24] E. Arikan and I. Telatar, "On the rate of channel polarization," in *Proc. IEEE Int. Symp. Info. Theory*, Jun. 2009, pp. 1493–1495.
[25] V. Strassen, "Asymptotische abschätzungen in Shannon's Informationstheorie," in *in Transactions of the 3rd Prague Conference on Information Theory, Statistical Decision Functions, Random Processes*. Prague: Pub. House of the Czechoslovak Academy of Sciences, 1962, pp. 689–723.
[26] S. Korada, E. Sasoglu, and R. Urbanke, "Polar codes: Characterization of exponent, bounds, and constructions," *IEEE Trans. Inf. Theory*, vol. 56, no. 12, pp. 6253–6264, Dec 2010.